

Intellectual Property (IP) in Untrusted Fabrication Environments

Stephen G. Baka & James T. Inge

Secure Computing and Communications Division

MacAulay-Brown, Inc.

4415 Pheasant Ridge Road, Suite 200, Roanoke, VA 24090

trust@macb.com

Abstract: *Manufacturing costs have driven many integrated circuit (IC) developers to a “fabless” model where IC fabrication is outsourced to untrusted foundries. In this paper we will present a threat model for this scenario, propose a concept of operations for re-establishing trust in ICs manufactured at untrusted foundries, and discuss results of a trust evaluation in an experimental scenario.*

Keywords: Trust; IC; ASIC; Risk Management; DoDI 5200.44; Intellectual Property (IP); Trusted Foundry; Fab; Trojan; Adversary; Threat Model

Introduction

Because of the expense of maintaining in-house fabrication capabilities, the modern *integrated circuit* (IC) development market necessitates a “fabless” structure for most IC development companies. A fabless development company (here referred to as *designer*) creates the original design, and then hands it off in an electronic format to a *fab*, a manufacturing facility that turns the design into a physical device. This source design is the *intellectual property* (IP) of the designer who desires it be protected from malicious modification.

The responsibilities of the fabless design house and the fab can vary considerably. A typical arrangement is for the designer to create the IC’s functional requirements, architectural schematics, and hardware description, usually in a *hardware description language* (HDL) such as VHDL or Verilog. The designer might also perform the first steps of synthesis and physical layout using an ASIC standard cell library chosen from among the process technologies supported by the fab. Prior to fab, the designer will also perform verification to ensure that the design they are handing off performs its intended function, as described by their product requirements. Thus, the *design* sent to the fab represents a *trusted*, functionally verified description of the intended product that synthesized into one of the fab’s supported ASIC standard cell libraries.

After the fab receives the trusted design, depending on its format, the fab may simply finish the synthesis, layout, mask generation, or IC fabrication. However, the eccentricities of modern semiconductor processes are making it increasingly more common for the fab to take additional steps. It may revisit each step of the designer’s process prior to the final fabrication steps in order to ensure that the designer’s implementation style is optimized for the selected process. Whether or not it revisits these steps, the fab has enough information about the design to insert new functionality. This is a concern when the designer is developing a device for use in a DoD system and the fab is not a trusted foundry.

Threat Model

When we claim *trust* in a device we are asserting that the fabricated device correctly implements all of the requirements defined by the designers and only those requirements, with no additional functionality. There are various threats to the proper fabrication of an ASIC after the trusted design information is handed over to the fab. Our threat model is predicated on some functional difference between the original trusted design and the final device produced by the fab. Any such difference is what we define as a *change*. When developing our change taxonomy we looked at the issue from two perspectives: intent and effect. These two perspectives delimit the start and end, respectively, of a change’s life-cycle.

Life-Cycle of a Change

All changes have an origin and start their life-cycle (**Figure 1**) in one of two ways: *intentionally* or *unintentionally*. Changes made with intent could come from an adversary, to undermine the design, or by the fab, to ensure proper yield and testability of the device. Thus, intentional changes fall into two subcategories – *malicious* changes, made by the adversary, and *benevolent* changes, made in good faith by the fab. Changes may also appear unintentionally, as an artifact of fabrication or Trojan insertion by an adversary.




Figure 1 – Life-Cycle of a Change

Any threat, regardless of original intent, may result in a *malignant*, *benign*, or *helpful* effect in the fabricated ASIC. This effect on the device or its parent system marks the end of the change's life-cycle. Note that the original intent of a change is not important when considering its actual effect on the fabricated device. Indeed, it may be impossible for a trust evaluator to distinguish a flaw inserted accidentally by the fab from an actual Trojan circuit. In fact, a clever adversary could disguise a Trojan to appear as a design flaw or synthesis artifact. Thus the labels *malignant*, *benign*, and *helpful* only apply to how a change affects the functionality, security, or reliability of the ASIC.

Solving the Trust in Fabrication Problem

As part of a two Phase SBIR program (OSD10-A08 [1]) funded by the US Army, MacB was tasked with addressing this threat model and establishing a concept of operations for re-establishing trust in ICs

manufactured at untrusted foundries. This effort culminated in a proof-of-concept experiment that demonstrated the efficacy of our threat model, *ASIC Change Detection Platform* (ASIC-CDP [2]), and trust evaluation methodology.

Concept of Operations

Our concept of operations (**Figure 2**) for establishing trust in ICs manufactured by untrusted foundries is based upon the following simplifying assumptions:

- trust in the ASIC *designer* and requirements
- access to a trusted source design or netlist
- no trust in the ASIC *fab* that constructs the device
- extraction of a gate-level netlist from the fabricated IC using proven delayering methods (e.g., a silicon-to-GDSII-to-Verilog netlist translation)
- trust in the netlist extraction entity [3] [4]

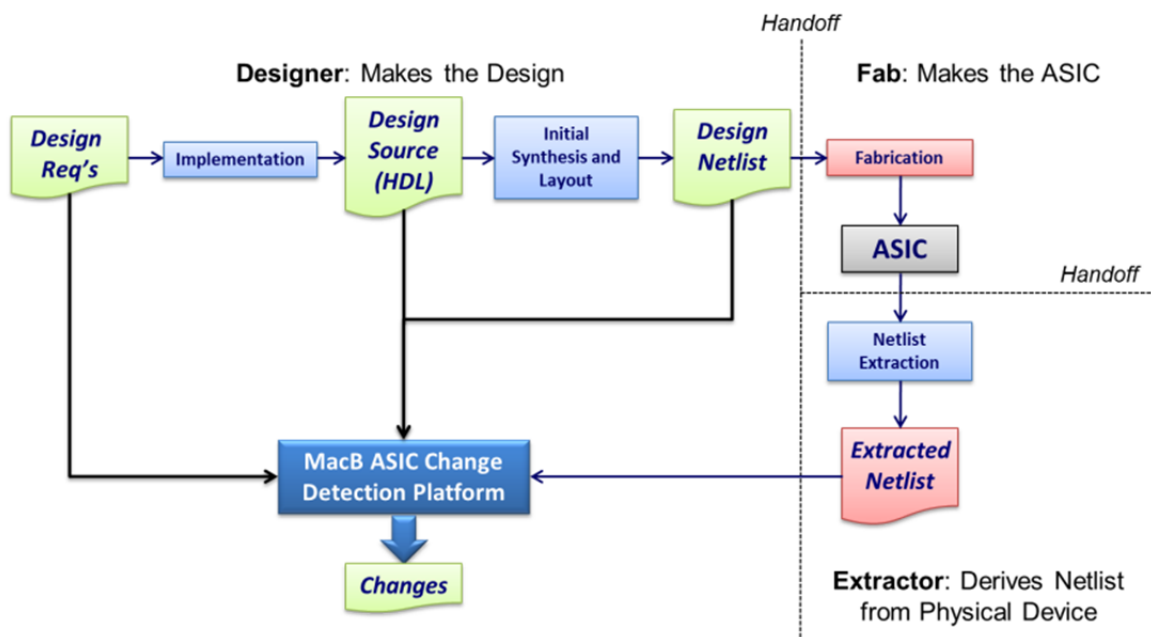


Figure 2 – Concept of Operations for Establishing Trust in IP Manufactured by Untrusted Foundry

Trust Evaluation Methods

In performing trust evaluations, we assumed the evaluators had full access to an array of industry-standard debug and verification tools, including schematic viewers, functional simulators, and formal methods. In addition to these, the evaluators were provided with custom trust evaluation methods via our ASIC-CDP. The platform included several methods for automated mapping, extraction, static detection, and transformation of digital circuits to rapidly isolate areas of interest.

Proof-of-Concept Evaluation

This effort culminated in a proof-of-concept experiment to determine the efficacy of our threat model, an *ASIC Change Detection Platform* (ASIC-CDP), and a trust evaluation methodology. As part of the experiment we formed a blue team (i.e., trusted design house or trust evaluator) and a red team (i.e., untrusted fab or adversary). The blue team created the representative design, setup its physical constraints, synthesized it for the targeted ASIC standard cell library, and handed off all trusted design data to the red team. Next, the red team applied our threat model to the representative design and produced several untrusted netlists for the blue team to evaluate. We assumed each netlist represented a design extracted by a trusted entity from a delayered and imaged sample device. Finally, the blue team applied MacB's CDP methods and trust evaluation techniques to the untrusted netlists and reported any detected changes. Throughout the process, the blue team was blinded to the changes inserted by the red team into the untrusted netlists.

Summary

In this paper we have outlined a concept-of-operations for how one might assure trust in an ASIC IP manufactured by an untrusted fabrication facility. By leveraging mature ASIC delayering and GDSII-to-gates translation capabilities one can recreate an EDA-compatible HDL netlist from the untrusted fabricated device. Trust evaluators, using our ASIC-CDP, can then identify suspicious design modifications. We demonstrated our trust technology and concept-of-operations to the US Army as part of the "Intellectual Property (IP) in Untrusted Fabrication Environments" Phase-I/II SBIR in which our

blue and red teams played out the roles of designer/evaluator and fab/adversary, respectively, across a series of test articles. We feel that our ASIC-CDP combined with our trust-by-evaluation process will play an important role in DOD's planning for implementation of trust policy as outlined in DODI 5200.44 [5].

Public Release Information

This paper is approved for public release; distribution is unlimited. Please contact MacAulay-Brown at trust@macb.com for presentation slides and additional details.

Acknowledgements

This material is based on work supported by OSD under Contract Number W31P4Q-11-C-0159, by the Defense Advanced Research Projects Agency (DARPA) under Contract Number HR0011-08-C-0007, and by the Army SBIR Programs office under Contract Number W31P4Q-12-C-0130. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of OSD, DARPA, or the Army SBIR Programs office.

References

- [1] OSD, "Intellectual Property (IP) in untrusted fabrication environments," DoD SBIR Resource Center, 2010. [Online]. Available: http://www.dodsbir.net/Sitis/archives_display_topic.asp?Bookmark=38366.
- [2] J. Graf, "Change Detection Platform for FPGA Trust," in *GOMAC*, Orlando, 2011.
- [3] Integra Technologies, "Military / Space / Aerospace Services," [Online]. Available: <http://www.integra-tech.com/Military-Space-Aerospace.html>. [Accessed December 2014].
- [4] Analytical Solutions, "Destructive Physical Analysis (DPA)," [Online]. Available: <http://www.asinm.com/services/dpa.aspx>. [Accessed December 2014].
- [5] DoD CIO/USD(AT&L), "DODI 5200.44 - Protection of Mission Critical Functions to Achieve Trusted Systems and Networks (TSN)," 2012.